arXiv:1702.07435v1 [cs.DS] 24 Feb 2017

Capacitated Center Problems with Two-Sided Bounds and Outliers

Hu Ding* Lunjia Hu† Lingxiao Huang† Jian Li†

February 27, 2017

Abstract

In recent years, the capacitated center problems have attracted a lot of research interest. Given a set of vertices V, we want to find a subset of vertices S, called centers, such that the maximum cluster radius is minimized. Moreover, each center in S should satisfy some capacity constraint, which could be an upper or lower bound on the number of vertices it can serve. Capacitated k-center problems with one-sided bounds (upper or lower) have been well studied in previous work, and a constant factor approximation was obtained.

We are the first to study the capacitated center problem with both capacity lower and upper bounds (with or without outliers). We assume each vertex has a uniform lower bound and a non-uniform upper bound. For the case of opening exactly k centers, we note that a generalization of a recent LP approach can achieve constant factor approximation algorithms for our problems. Our main contribution is a simple combinatorial algorithm for the case where there is no cardinality constraint on the number of open centers. Our combinatorial algorithm is simpler and achieves better constant approximation factor compared to the LP approach.

*Computer Science and Engineering, Michigan State University, East Lansing, MI, USA

†Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China

1 introduction

The k-center clustering is a fundamental problem in theoretical computer science and has numerous applications in a variety of fields. Roughly speaking, given a metric space containing a set of vertices, the k-center problem asks for a subset of k vertices, called centers, such that the maximum radius of the induced k clusters is minimized. Actually k-center clustering falls in the umbrella of the general *facility location* problems which have been extensively studied in the past decades. Many operation and management problems can be modeled as facility location problems, and usually the input vertices and selected centers are also called "clients" and "facilities" respectively. In this paper, we consider a significant generalization of the k-center problem, where each vertex is associated with a capacity interval; that is, the cardinality of the resulting cluster centered at the vertex should satisfy the given lower and upper capacity bounds (the formal definition is shown in Section 1.2). In addition, we also consider the case where a given number of vertices may be excluded as outliers.

Besides being a natural combinatorial problem on its own, the k-center problem with both capacity upper and lower bounds is also strongly motivated by several realistic issues raised in a variety of application contexts.

1. In the context of facility location, each open facility may be constrained by the maximum number of clients it can serve. The capacity lower bounds also come naturally, since an open facility needs to serve at least a certain number of clients in order to generate profit.

2. Several variants of the k-center clustering have been used in the context of preserving privacy in publication of sensitive data (see e.g., (Aggarwal *et al.* , 2010; Li *et al.* , 2010; Sweeney, 2002)). In such applications, it is important to have an appropriate lower bound for the cluster sizes, in order to protect the privacy to certain extent (roughly speaking, it would be relatively easier for an adversary to identify the clients inside a too small cluster).

3. Consider the scenario where the data is distributed over the nodes in a large network. We would like to choose k nodes as central servers, and aggregate the information of the entire network. We need to minimize the delay (i.e., minimize the cluster radius), and at the same time consider the balancedness, for the obvious reason that the machines receiving too much data could be the bottleneck of the system and the ones receiving too little data is not sufficiently energy-efficient (Dick *et al.* , 2015).

Our problem generalizes the classic k-center problem as well as many important variants studied by previous authors. The optimal approximation results for the classic k-center problem appeared in the 80's: Gonzalez (1985) and Hochbaum & Shmoys (1985) provided a 2-approximation in a metric graph; moreover, they proved that any approximation ratio $c < 2$ would imply $P = NP$. The first study on capacitated (with only upper bounds) k-center clustering is due to Barilan *et al.* (1993) who provided a 10-approximation algorithm for uniform capacities (i.e., all the upper bounds are identical). Further, Khuller & Sussmann (2000) improved the approximation ratio to be 6 and 5 for hard and soft uniform capacities, respectively. * The recent breakthrough for non-uniform (upper) capacities is due to Cygan *et al.* (2012). They developed the first constant approximation algorithm based on LP rounding, though their approximation ratio is about hundreds. Following

* We can open more than one copies of a facility in the same node in the soft capacity version. But in the hard capacity version, we can only open at most one copy.

this work, An *et al.* (2015) provided an approximation algorithm with the much lower approximation ratio 9. On the imapproximability side, it is impossible to achieve an approximation ratio lower than 3 for non-uniform capacities unless $P = NP$ (Cygan *et al.* , 2012).

For the ordinary k-center with outliers, a 3-approximation algorithm was obtained by Charikar *et al.* (2001). Kociumaka & Cygan (2014) studied k-center with non-uniform upper capacities and outliers, and provided a 25-approximation algorithm.

k-center clustering with lower bounds on cluster sizes was first studied in the context of privacy-preserving data management (Sweeney, 2002). Aggarwal *et al.* (2010) provided a 2-approximation and a 4-approximation for the cases without and with outliers, respectively. Further, Ene *et al.* (2013) presented a near linear time $(4 + \epsilon)$-approximation algorithm in constant dimensional Euclidean space. Note that both (Aggarwal *et al.* , 2010; Ene *et al.* , 2013) are only for uniform lower bounds. Recently, Ahmadian & Swamy (2016) provided a 3-approximation and a 5-approximation for the non-uniform lower bound case without and with outliers.

Our main results. To the best of our knowledge, we are the first to study the k-center with both capacity lower and upper bounds (with or without outliers). Given a set V of n vertices, we focus on the case where the capacity of each vertex $u \in V$ has a uniform lower bound $L_u = L$ and a non-uniform upper bound U_u. Sometimes, we consider a generalized supplier version where we are only allowed to open centers among a facility set $\mathcal{F}$, see Definition 1 for details. We mainly provide first constant factor approximation algorithms for the following variants, see Table 1 for other results.

1. (L,U,soft-$\emptyset$,p)-Center (Section 2.2): In this problem, both the lower bounds and the upper bounds are uniform, i.e., $L_u = L, U_u = U$ for all $u \in V$. The number of open centers can be arbitrary, i.e., there is no requirement to choose exactly k open centers. Moreover, we allow multiple open centers at a single vertex $u \in V$ (i.e., soft capacity). We may exclude $n - p$ outliers. We provide the first polynomial time combinatorial algorithm which can achieve an approximate factor of 5.

2. (L,$\{U_u\}$,$\emptyset$,p)-Center(Section 2.3): In this problem, the lower bounds are uniform, i.e., $L_u = L$ for all $u \in V$, but the upper bound can be nonuniform. The number of open centers can be arbitrary. We may exclude $n-p$ outliers. We provide the first polynomial time combinatorial 11-approximation for this problem.

3. (L,$\{U_u\}$,k)-Center (Section 3.3): In this problem, we would like to open exactly k centers, such that the maximum cluster radius is minimized. All vertices have the same capacity lower bounds, i.e., $L_u = L$ for all $u \in V$. But the capacity upper bounds may be nonuniform, i.e., each vertex u has an individual capacity upper bound U_u. Moreover, we do not exclude any outlier. We provide the first polynomial time 9-approximation algorithm for this problem, based on LP rounding.

4. (L,$\{U_u\}$,k,p)-Center (Section 3.3): This problem is the outlier version of the (L,$\{U_u\}$,k)-Center problem. The problem setting is exactly the same except that we can exclude $n - p$ vertices as outliers. We provide a polynomial time 25-approximation algorithm for this problem.

Our main techniques. In Section 2, we consider the first two variants which allow to open arbitrarily many centers. We design simple and faster combinatorial algorithms which can achieve

Problem Setting		Approximation Ratio	
		Center Version	Supplier Version
Without k Constraint	$(L,U,\text{soft-}\emptyset,p)$	5	5
	$(L,U,\emptyset,p)$	10	23
	$(L,\{U_u\},\text{soft-}\emptyset,p)$	11	11
	$(L,\{U_u\},\emptyset,p)$	11	25
With k Constraint	(L,U,k)	6	9
	$(L,\{U_u\},k)$	9	13
	$(L,U,\text{soft-}k,p)$	13	13
	(L,U,k,p)	23	23
	$(L,\{U_u\},\text{soft-}k,p)$	25	25
	$(L,\{U_u\},k,p)$	25	25

Table 1: A summarization table for our results in this paper.

better constant approximation ratios compared to the LP approach. For the simpler case $(L,U,\text{soft-}\emptyset,p)$-Center, we construct a data structure for all possible open centers. We call it a *core-center tree* (CCT). Our greedy algorithm mainly contains two procedures. The first procedure *pass-up* greedily assigns vertices to open centers from the leaves of CCT to the root. After this procedure, there may exist some unassigned vertices around the root. We then introduce the second procedure called *pass-down*, which assigns these vertices in order by finding an *exchange route* each time. For the more general case $(L,\{U_u\},\emptyset,p)$-Center, our greedy algorithm is similar but somewhat more subtle. We still construct a CCT and run the pass-up procedure. Then we obtain an open center set F, which may contain redundant centers. However, since we deal with hard capacities and outliers, we need to find a non-redundant open center set which is not 'too far' from F (see Section 2.3 for details) and have enough total capacities. Then by a pass-down procedure, we can assign enough vertices to their nearby open centers.

In Section 3 and 3.3, we consider the last two variants which require to open exactly k centers. We generalized the LP approach developed for k-center with only capacity upper bounds (An *et al.* , 2015; Kociumaka & Cygan, 2014) and obtain constant approximation schemes for two-sided capacitated bounds. Due to the lack of space, we defer many details and proofs to a full version.

1.1 Other Related Work

The classic k-center problem is quite fundamental and has been generalized in many ways, to incorporate various constraints motivated by different application scenarios. Recently, Fernandes *et al.* (2016) also provided constant approximations for the fault-tolerant capacitated k-center clustering. Chen *et al.* (2016) studied the matroid center problem where the selected centers must form an independent set of a given matroid, and provided constant factor approximation algorithms (with or without outliers).

There is a large body of work on approximation algorithms for the facility location and k-median problems (see e.g., (Arya *et al.* , 2004; Charikar & Guha, 2005; Charikar *et al.* , 1999; Guha & Khuller, 1999; Jain *et al.* , 2002; Jain & Vazirani, 2001; Korupolu *et al.* , 2000; Li, 2013; Li & Svensson, 2016)). Moreover, Dick *et al.* (2015) studied multiple balanced clustering prob-

lems with uniform capacity intervals, that is, all the lower (upper) bounds are identical; they also consider the problems under the stability assumption.

1.2 Preliminaries

In this paper, we usually work with the following more general problem, called the capacitated k-supplier problem. It is easy to see it generalizes the capacitated k-center problem. The formal definition is as follows.

Definition 1. (Capacitated k-supplier with two-sided bounds and outliers) Suppose that we have

1. Two integers $k, p \in \mathbb{Z}_{\geq 0}$;
2. A finite set $\mathcal{C}$ of clients, and a finite set $\mathcal{F}$ of facilities;
3. A symmetric distance function $\mathrm{d} : (\mathcal{C} \cup \mathcal{F}) \times (\mathcal{C} \cup \mathcal{F}) \to \mathbb{R}_{\geq 0}$ satisfying the triangle inequality;
4. A capacity interval $[L_u, U_u]$ for each facility $u \in \mathcal{F}$, where $L_u, U_u \in \mathbb{Z}_{\geq 0}$ and $L_u \leq U_u$.

Our goal is to find a client set $C \subseteq \mathcal{C}$ of size at least p, an open facility set $F \subseteq \mathcal{F}$ of size exactly k, and a function $\phi : C \to F$ satisfying that $L_u \leq |\phi^{-1}(u)| \leq U_u$ for each $u \in F$, which minimize the maximum cluster radius $\max_{v \in C} \mathrm{d}(v, \phi(v))$. If the maximum cluster radius is at most r, we call the tuple (C, F, ϕ) a distance-r solution.

We denote the above problem as $(\{L_u\},\{U_u\},k,p)$-Supplier. If the lower bounds are uniform ($L_u = L$ for all $u \in \mathcal{F}$), we use L in place $\{L_u\}$, e.g., $(L,\{U_u\},k)$-Supplier. Similarly, if the upper bounds are uniform ($U_u = U$ for all $u \in \mathcal{F}$), we use U in place $\{U_u\}$. If there is no constraint to open k centers, we use $\emptyset$ to replace k, e.g., $(L,\{U_u\},\emptyset,p)$-Supplier. Also note that the capacitated k-center problem with two-sided bounds and outliers is a special case by letting $V = \mathcal{C} = \mathcal{F}$, we denote it the $(\{L_u\},\{U_u\},k,p)$-Center problem.

By the similar approach of Kociumaka & Cygan (2014), we can reduce the $(\{L_u\},\{U_u\},k,p)$-Supplier problem to a simpler case. We first introduce some definitions.

Definition 2. (Induced distance function) We say the distance function $\mathrm{d}_G : (\mathcal{C} \cup \mathcal{F}) \times (\mathcal{C} \cup \mathcal{F}) \to \mathbb{R}_{\geq 0}$ is induced by an undirected unweighted connected graph $G = (\mathcal{C} \cup \mathcal{F}, E)$ if

1. $\forall (u, v) \in E$, we have $u \in \mathcal{F}$ and $v \in \mathcal{C}$.
2. $\forall a_1, a_2 \in \mathcal{C} \cup \mathcal{F}$, the distance $\mathrm{d}_G(a_1, a_2)$ between a_1 and a_2 equals to the length of the shortest path from a_1 to a_2.

Definition 3. (Induced $(\{L_u\},\{U_u\},k,p)$-Supplier instance) An $(\{L_u\},\{U_u\},k,p)$-Supplier instance is called an induced $(\{L_u\},\{U_u\},k,p)$-Supplier instance if the following properties are satisfied:

1. The distance function d_G is induced by an undirected connected graph $G = (\mathcal{C} \cup \mathcal{F}, E)$.
2. The optimal capacitated k-supplier value is at most 1.

Moreover, we say this instance is induced by G.

When the graph of interest G is clear from the context, we will use d instead of d_G for convenience. We then show a reduction from solving the generalized $(\{L_u\},\{U_u\},k,p)$-Supplier problem to solving induced $(\{L_u\},\{U_u\},k,p)$-Supplier instances by Lemma 4. The proof can be found in Appendix A.

Lemma 4. Suppose we have a polynomial time algorithm A that takes as input any induced $(\{L_u\},\{U_u\},k,p)$-Supplier instance, and outputs a distance-ρ solution. Then, there exists a ρ-approximation algorithm for the $(\{L_u\},\{U_u\},k,p)$-Supplier problem with polynomial running time.

By Lemma 4, we focus on designing an algorithm A for different variants of the induced $(\{L_u\},\{U_u\},k,p)$-Supplier instances.

2 Capacitated Center with Two-Sided Bounds and Outliers

In this section, we consider the version that the number of open centers can be arbitrary. By the LP approach in Section 3.3 and enumerating the number of open centers, we can achieve approximation algorithms for different variants in this case. However, the approximation factor is not small enough. In this section, we introduce a new greedy approach in order to achieve better approximation factors. Since our algorithm is combinatorial, it is easier to be implemented and saves the running time compared to the LP approach.

2.1 Core-center tree (CCT)

Consider the $(L,\{U_u\},\emptyset,p)$-Supplier problem. By Lemma 4, we only need to consider induced $(L,\{U_u\},\emptyset,p)$-Supplier instances induced by an undirected unweighted connected graph $G = (\mathcal{C} \cup \mathcal{F}, E)$. We first propose a new data structure called *core-center tree (*CCT*)* as follows.

Definition 5. (Core-center tree (CCT)) Given an induced $(L,\{U_u\},\emptyset,p)$-Supplier instance induced by an undirected unweighted connected graph $G = (\mathcal{C} \cup \mathcal{F}, E)$, we call a tree $T = (\mathcal{F}, E_T)$ a core-center tree(CCT) if the following properties hold.

1. For each edge $(u, u') \in E_T$, we have $\mathrm{d}_G(u, u') \leq 2$;
2. Suppose the root of T is at layer 0. Denote I to be the set of vertices in the even layers of T. We call I the core-center set of T. For any two distinct vertices $u, u' \in I$, we have $\mathrm{d}_G(u, u') \geq 3$.

Lemma 6. Given an induced $(L,\{U_u\},\emptyset,p)$-Supplier instance induced by an undirected unweighted connected graph $G = (\mathcal{C} \cup \mathcal{F}, E)$, we can construct a CCT in polynomial time.

Proof. We first construct a graph G^2 on $\mathcal{F}$ as follows: for each pair u_1, u_2 in $\mathcal{F}$ with distance at most 2, we add an edge (u_1, u_2) in G^2. Observe that G^2 is connected by Definition 2. We then construct a spanning tree T of G^2 satisfying that all facilities in even layers form an independent set of G^2. It is not hard to verify that such a tree is a CCT. We build T as follows. The above property directly holds from our construction.

1. Initially, we randomly pick a facility $u \in \mathcal{F}$ as the root of T. We then pick all adjacent facilities of u in G^2 as its children (layer 1).

2. By a modified BFS, we continue to construct layer 2 and layer 3. Each time we pick a facility w in layer 1. We iteratively pick an adjacent facility w' of w in G^2 which has not been scanned as a child of w. After we append w' to layer 2, we immediately pick all unscanned neighbors of w' in G^2 as the children of w' (append them in layer 3).

3. We then iteratively construct T until all facilities in $\mathcal{F}$ have been scanned. Each iteration, we build two consecutive layers: an odd layer and an even layer.

□

For any $u \in \mathcal{F}$, denote $N_G[u] = \{v \in \mathcal{C} : (u, v) \in E\}$ to be the collection of all neighbors of $u \in \mathcal{F}$. [†] W.l.o.g., we assume that $U_u \leq |N_G(u)|$ for every facility $u \in \mathcal{F}$ in this section. In fact, we can directly delete all $u \in \mathcal{F}$ satisfying that $|N_G[u]| < L$ from the facility set $\mathcal{F}$, since u can not be open in any optimal feasible solution. [‡] Otherwise if $L \leq |N_G[u]| < U_u$, we set $U_u \leftarrow \min\{U_u, |N_G[u]|\}$, which has no influence on any optimal feasible solution of the induced $(L,\{U_u\},\emptyset,p)$-Supplier instance. The following lemma gives a useful property of CCT.

Lemma 7. Given an induced $(L,\{U_u\},\emptyset,p)$-Supplier instance induced by an undirected unweighted connect graph $G = (\mathcal{C} \cup \mathcal{F}, E)$, and a core-center tree $T = (\mathcal{F}, E_T)$, suppose I is the core-center set of T. Then, we can construct a function $\xi : \mathcal{C} \to \mathcal{F}$ satisfying the following properties in polynomial time.

1. For all $v \in \mathcal{C}$, we have $(\xi(v), v) \in E$;

2. For all $u \in I$, we have $|\xi^{-1}(u)| \geq L$.

Proof. Firstly, for each pair $u \in I$ and $v \in N_G[u]$, we define $\xi(v) = u$. We can make this mapping since for each pair $u_1, u_2 \in I$, we have $N_G[u_1] \cap N_G[u_2] = \emptyset$ by Definition 5. For the rest clients $v \in \mathcal{C}$, we define $\xi(v)$ to be an arbitrary facility $u \in \mathcal{F}$ adjacent to v.

By the above construction, the first constraint is satisfied naturally. The second constraint is satisfied by the fact that $|\xi^{-1}(u)| \geq |N_G[u]| \geq L$ for all $u \in I$. □

2.2 A Simple Case: (L,U,soft-$\emptyset$,p)-Supplier

We first consider a simple case where the capacity bounds (upper and lower) are uniform and soft. In this setting, we want to find an open facility set $F = \{u_i \mid u_i \in \mathcal{F}\}_i$. Note that we allow multiple open centers in F. We also need to find an assignment function $\phi : \mathcal{C} \to F$, representing that we assign every client $v \in \mathcal{C}$ to facility $\phi(v)$. The main theorem is as follows.

Theorem 8. (main theorem) There exists a 5-approximation polynomial time algorithm for the $(L,U$,soft-$\emptyset,p)$-Supplier problem.

By Lemma 4, we only consider induced $(L,U$,soft-$\emptyset,p)$-Supplier instances. Given an induced $(L,U$,soft-$\emptyset,p)$-Supplier instance induced by an undirected unweighted connect graph $G = (\mathcal{C} \cup$

[†] If $u \in \mathcal{C}$ is also a client, then $u \in N_G[u]$.

[‡] If this deletion causes the induced graph unconnected, similar to Lemma 6 in (Kociumaka & Cygan, 2014), we divide the graph into different connected components, and consider each smaller induced instance based on different connected components.

$\mathcal{F}, E)$, recall that we can assume $|N_G[u]| \geq U_u \geq L$ for each $u \in \mathcal{F}$. We first construct a CCT $T = (\mathcal{F}, E_T)$ rooted at node u^*, and a function $\xi : \mathcal{C} \to \mathcal{F}$ satisfying Lemma 7. For a facility set $P \subseteq \mathcal{F}$, we denote $\xi^{-1}(P) = \bigcup_{u \in P} \xi^{-1}(u)$ to be the collection of clients assigning to some facility in P by ξ.

Our algorithm mainly includes two procedures. The first procedure is called *pass-up*, which is a greedy algorithm to map clients to facilities from the leaves of T to the root. After the 'pass-up' procedure, we still leave some unassigned clients nearby the root. Then we use a procedure called *pass-down* to allocate those unassigned clients by iteratively finding an *exchange route*. In the following, we give the details of both procedures.

Procedure Pass-Up. Assume that $|\mathcal{C}| = aL + b$ for some $a \in \mathbb{N}$ and $0 \leq b \leq L - 1$. In this procedure, we will find an open facility set F of size a. We also find an assignment function ϕ which assigns aL clients to some nearby facility in F except a client set $S \subseteq \mathcal{C}$. Here, S is a collection of b clients in $\xi^{-1}(u^*)$ nearby the root u^*. Our main idea is to open facility centers from the leaves of CCT T to the root iteratively. During opening centers, we assign exactly L 'close' clients to each center. This is the reason that there are b unassigned clients after the whole procedure.

We then describe an iteration of pass-up. Assume that I is the core-center set of T. At the beginning, we find a non-leaf vertex $u \in I$ satisfying that all of its grandchildren (if exists) are leaves. We denote $P \subseteq \mathcal{F}$ to be the collection of all children and all grandchildren of u. In the next step, we consider all unscanned clients in $\xi^{-1}(P)$, [§] and assign them to the facility u. We want that each center at u serves exactly L centers. However, there may exist one center at u serving less than L unscanned clients in $\xi^{-1}(P)$. We assign some clients in $\xi^{-1}(u)$ to this center such that it also serves exactly L clients. After this iteration, we delete the subtree rooted at u from T except u itself.

Finally, the root u^* will become the only remaining node in T. We open multiple centers at u^*, each serving exactly L clients in $\xi^{-1}(u^*)$, until there are less than L unassigned clients. See Algorithm 1 for details. We have the following lemma.

Lemma 9. Given an induced (L,U,soft-$\emptyset$,p)-Supplier instance induced by an undirected unweighted connect graph $G = (\mathcal{C} \cup \mathcal{F}, E)$, assume that $|\mathcal{C}| = aL + b$ for some $a \in \mathbb{N}$ and $0 \leq b \leq L - 1$. The output of Algorithm 1 satisfies the following properties:

1. Each open facility $u_j \in F$ satisfies that $u_j \in I$, and $|F| = a$;
2. The unassigned client set $S \subseteq \xi^{-1}(u^*)$, and $|S| = b$;
3. For each facility $u_i \in F$, we have $|\phi^{-1}(u_i)| = L$.
4. For each client $v \in \mathcal{C} \setminus S$, $\phi(v)$ is either $\xi(v)$, or the parent of $\xi(v)$ in T, or the grandparent of $\xi(v)$ in T. Moreover, we have $\mathrm{d}_G(v, \phi(v)) \leq 5$.

Proof. We first prove the feasibility of Algorithm 1. The feasibility of Line 7 follows from the fact that $|\xi^{-1}(u)| \geq L$ by Lemma 7. Since $0 \leq q \leq L - 1$ by Line 6, we can always pick $L - q$ clients from $\xi^{-1}(u)$. The feasibility of Line 9 follows from the fact that $|A| = l + L - q = (t + 1)L$. Since we open $(t + 1)$ centers at u, it is able to assign exactly L clients in A to each center.

[§]Here, unscanned clients are those clients that have not been assigned by ϕ before this iteration.

Algorithm 1: Pass-Up

Input: an induced (L,U,soft-$\emptyset$,p)-SUPPLIER instance induced by $G = (\mathcal{C} \cup \mathcal{F}, E)$, a CCT $T = (\mathcal{F}, E_T)$, and a function $\xi : \mathcal{C} \to \mathcal{F}$;

Initialize $S \leftarrow \mathcal{C}$, $T' \leftarrow T$, $j \leftarrow 0$;

while $u^* \in T'$ **do**

  If the root u^* is the only node of T', we let $u \leftarrow u^*$. Otherwise, arbitrarily pick a non-leaf vertex $u \in I$ in T' whose all grandchildren (if exists) are leaves of T' ;

  Denote the subtree of T' rooted at u by $\hat{T}$. Denote P to be the collection of all facilities in $\hat{T} \setminus \{u\}$;

  Let $l \leftarrow |\xi^{-1}(P) \cap S|$. Assume that $l = tL + q$ for some $t \in \mathbb{N}$ and $0 \leq q \leq L - 1$;

  Arbitrarily pick $L - q$ clients from $\xi^{-1}(u)$ to form a set K. Let $A \leftarrow (\xi^{-1}(P) \cap S) \cup K$;

  Let $u_{j+1} \leftarrow u, u_{j+2} \leftarrow u, \cdots, u_{j+t+1} \leftarrow u$;

  For each center u_{j+i} $(1 \leq i \leq t+1)$, assign exactly L clients $v \in A$ to u_{j+i}, i.e., let $\phi(v) = u_{j+i}$;

  Let $j \leftarrow j + t + 1$, $S \leftarrow S \setminus A$, $T' \leftarrow (T' \setminus \hat{T}) \cup \{u\}$;

  if $u = u^*$ **then**

    Let $l' = |\xi^{-1}(u) \cap S|$. Assume that $l' = t'L + q'$ for some $t \in \mathbb{N}$ and $0 \leq q \leq L - 1$;

    Arbitrarily pick $t'L$ clients from $\xi^{-1}(u) \cap S$ to form a set K';

    Let $u_{j+1} \leftarrow u, u_{j+2} \leftarrow u, \cdots, u_{j+t'} \leftarrow u$;

    For each center u_{j+i} $(1 \leq i \leq t')$, assign exactly L clients $v \in K'$ to u_{j+i}, i.e., let $\phi(v) = u_{j+i}$;

    Let $j \leftarrow j + t'$, $S \leftarrow S \setminus K'$, $T' \leftarrow \emptyset$;

Output: $F = \{u_1, u_2, \cdots, u_j\}$, $\phi : (\mathcal{C} \setminus S) \to F$ and S.

Then we prove the properties of the output. The first three properties mainly follow from the fact that $|\mathcal{C}| = aL + b$ and we assign exactly L clients to each open center. We only need to verify that $S \subseteq \xi^{-1}(u^*)$. By Line 3, we always pick $u = u^*$ in the last iteration of Algorithm 1. By Line 11-16, this fact is obvious. For each center at u, it only serves L clients in $(\xi^{-1}(P) \cap S) \cup K$. By the definition of P and K, we conclude the first part of the last property. Moreover, we have $\mathrm{d}_G(v, \xi(v)) = 1$ by Lemma 7 and $\mathrm{d}_G(\xi(v), u_{\phi(v)}) \leq 4$ by Definition 5. By the triangle inequality, we have $\mathrm{d}_G(v, \phi(v)) \leq 5$. □

Procedure Pass-Down. After the procedure *pass-up*, we still leave an unassigned client set S of size b. However, our goal is to serve at least p clients. Therefore, we need to modify the assignment function ϕ and serve more clients.

The procedure pass-down handles the remaining b clients in S one by one, see Algorithm 2 for details. At the beginning of pass-down, we initialize an 'unscanned' client set $B \leftarrow \mathcal{C} \setminus S$, i.e., B is the collection of those clients allowing to be reassigned by pass-down. In each iteration, we arbitrarily pick a client $v \in S$ and assign it to the root node u^*. However, if each open facility at u^* has already served U_{u^*} clients by ϕ, assigning v to u^* will violate the capacity upper bound. In this case, we actually find an open center $u_j \in F$ such that $|\phi^{-1}(u_j)| < U_j$, i.e., there are less than U_j clients assigned to u_j by ϕ. We then construct an *exchange route* consisting of open facilities in F. We first find a sequence of nodes $w_0 = u^*, w_1, \cdots, w_m = u_j$ in T satisfying that w_i is the grandparent of w_{i+1} in the core-center tree T for all $0 \leq i \leq m-1$. Then for each node w_i $(1 \leq i \leq m-1)$, we pick a client $v_i \in \xi^{-1}(w_i)$ which has not been reassigned so far. We call such a sequence of clients $v, v_1, \ldots, v_{m-1}$ an exchange route. Our algorithm is as follows: 1) we assign v to $\phi(v_1)$; 2) we iteratively reassign v_i to $\phi(v_{i+1})$ in order $(1 \leq i \leq m-2)$; 3) finally we reassign v_{m-1} to u_j. We then mark all clients v_i $(1 \leq i \leq m-1)$ in the exchange route by removing them from the 'unscanned' client set B. Note that our exchange route only increases the number of clients assigned to u_j by one. We will prove such an exchange route always exists in each iteration. Thus in each iteration, the procedure pass-down assigns one more client $v \in S$ to some open facility in F. We will argue that there are at least p clients served by F at the end of pass-down.

Now we prove the following lemma. Note that Theorem 8 can be directly obtained by Lemma 4 and Lemma 10.

Lemma 10. Algorithm 2 outputs a distance-5 solution (C, F, ϕ) of the given induced (L,U,soft-$\emptyset$,p)-Supplier instance induced by $G = (\mathcal{C} \cup \mathcal{F}, E)$ in polynomial time.

Proof. We first verify the feasibility of Algorithm 2. The feasibility of Line 8 follows from the fact that $u^*, u_j \in I$ by Lemma 9. Then we only need to show that an exchange route $\{v_0, \cdots, v_{m-1}\}$ described in Line 9-14 must exist in each iteration. On one hand, we verify that an exchange route $v, v_1, \ldots, v_{m-1}$ in Line 9 always exists. Since we assign one more client in S in each iteration. Thus, there are at most b iterations by Property 2 in Lemma 9. Then the node $w_t \in I$ appears at most b times in the sequence in Line 8, and at most $b < L$ clients in $\xi^{-1}(w_t)$ are removed from B at the end of the algorithm. By Lemma 7, we also have $|\xi^{-1}(w_t)| \geq L$. Thus we always have $\xi^{-1}(w_t) \cap B \neq \emptyset$ in Line 12, which proves the existence of an exchange route. On the other hand, for each open facility $\phi(v_i)$ $(1 \leq i \leq m-1)$, the number of clients served by $\phi(v_i)$ dose not change after Line 12. It is because we reassign v_{i-1} to $\phi(v_i)$, and remove v_i from it.

We then show that for all $v \in \mathcal{C}$, we have $\mathrm{d}_G(v, \phi(v)) \leq 5$ at the end of the algorithm. In Line

Algorithm 2: Pass-Down

Input: an induced (L,U,soft-$\emptyset$,p)-Supplier instance induced by $G = (\mathcal{C} \cup \mathcal{F}, E)$, a CCT $T = (\mathcal{F}, E_T)$, a function $\xi : \mathcal{C} \to \mathcal{F}$, an open facility set $F = \{u_1, u_2, \cdots, u_a\}$, an unassigned client set S, and a function $\phi : (\mathcal{C} \setminus S) \to F$;
Initialize $B \leftarrow \mathcal{C} \backslash S$;
while $S \neq \emptyset$ *and* $\exists 1 \leq j \leq a$, $|\phi^{-1}(j)| < U_j$ **do**
    Arbitrarily pick a client $v \in S$ and an open facility u_j $(1 \leq j \leq a)$ satisfying that $|\phi^{-1}(u_j)| < U_j$;
    if $u_j = u^*$ **then**
        Let $\phi(v) \leftarrow u_j$, $S \leftarrow S \setminus \{v\}$;
    else
        Let $w_0 = u^*, w_1, \cdots, w_m = u_j$ be the sequence of nodes in T where w_i is the grandparent of w_{i+1} for all $1 \leq i \leq m-1$.
        Let $v_0 \leftarrow v$. For every $1 \leq i \leq m-1$, arbitrarily pick a client $v_i \in \xi^{-1}(w_i) \cap B$;
        for $i = 0, \cdots, m-2$ **do**
            Reassign $\phi(v_i) \leftarrow \phi(v_{i+1})$;
        Reassign $\phi(v_{m-1}) \leftarrow u_j$;
        Let $S \leftarrow S \setminus \{v\}$, $B \leftarrow B \setminus \{v_1, v_2, \cdots, v_{m-1}\}$;
        Let $\phi(v_{m-1}) \leftarrow u_j$, $S \leftarrow S \setminus \{v\}$;
Output: $C \leftarrow \mathcal{C} \setminus S$, $F = \{u_1, u_2, \cdots, u_a\}$ and $\phi : C \to F$;

11, we reassign a client $v_{t-1} \in \xi^{-1}(w_{t-1})$ to $\phi(v_t)$. Since w_t is the grandchild of w_{t-1}, we have $\mathrm{d}_G(v_t, w_{t-1}) \leq 4$ by Property 4 in Lemma 9. Combining with $v_{t-1} \in \xi^{-1}(w_{t-1})$, we conclude that $\mathrm{d}_G(v_{t-1}, \phi(w_t)) \leq 5$. Then by Property 4 in Lemma 9, we finish the proof.

Finally, we show that the output client set C is of size at least p. In fact, we only need to prove that $\sum_{j=1}^{a} U \geq p$, i.e., $aU \geq p$. Since the number of open facility centers in the optimal solution served at least p clients is at most $a = \lfloor |\mathcal{C}|/L \rfloor$, we have $p \leq aU$. □

(L,U,$\emptyset$,p)-Center. Consider the (L,U,$\emptyset$,p)-Center problem with hard capacities. We first treat a given induced (L,U,$\emptyset$,p)-Center instance as an induced (L,U,soft-$\emptyset$,p)-Center instance. Then we apply Theorem 8 and obtain a 5-approximation solution (C, F, ϕ). Since the two instances are induced by the same connected graph, the optimal capacitated center value of the induced (L,U,soft-$\emptyset$,p)-Center instance is at most the optimal capacitated center value 1 of the induced (L,U,$\emptyset$,p)-Center instance. Therefore, we know $\max_{v \in C} \mathrm{d}(v, \phi(v)) \leq 5$. Since we have hard capacities, we still need to modify F to be a single set. In fact, we can choose arbitrary vertex $v_i \in \phi^{-1}(u_i)$ to replace each $u_i \in F$ as an open center, and assign all vertices in $\phi^{-1}(u_i)$ to v_i. Note that the distance between any $v \in \phi^{-1}(u_i)$ and v_i is at most 5 and the new u_i is at most 10. Thus we have the following theorem.

Theorem 11. There exists a 10-approximation polynomial time algorithm for the (L,U,$\emptyset$,p)-Center problem.

2.3 $(L,\{U_u\},\emptyset,p)$-Center

In this subsection, we consider a more complicated case where the capacity upper bounds are non-uniform, and each vertex has a hard capacity. Our main theorem is as follows.

Theorem 12. (main theorem) There exists an 11-approximation polynomial time algorithm for the $(L,\{U_u\},\emptyset,p)$-Center problem.

By Lemma 4, we only need to consider induced $(L,\{U_u\},\emptyset,p)$-Supplier instances. For an induced $(L,\{U_u\},\emptyset,p)$-Supplier instance induced by an undirected unweighted connected graph $G=(V=\mathcal{C}\cup\mathcal{F},E)$, recall that we can assume $U_u \le |N_G(u)|$ for every vertex $u\in\mathcal{F}$. [¶] Since we consider the center version, every vertex $v\in\mathcal{C}$ has an individual capacity interval $[L,U_v]$ and can be opened as a center as well. This fact is useful for our following algorithm and is the reason why we do not consider the supplier version in this subsection.

Similar to $(L,U,\text{soft-}\emptyset,p)$-Center, our algorithm first computes a core-center tree $T=(\mathcal{F},E)$ rooted at u^*, a core-center set I and a function ξ described as in Lemma 7. Assume that $|\mathcal{C}|=aL+b$ for some $a\in\mathbb{N}$ and $0\le b\le L-1$. Note that the procedure *pass-up* algorithm does not depend on the capacity upper bounds. Therefore, we still use the procedure *pass-up* to compute an open set $F=\{u_1,u_2,\cdots,u_a\}$, an unassigned set $S\subseteq\xi^{-1}(u^*)$ of size $b<L$, and a function $\phi:(\mathcal{C}\setminus S)\to F$.

However, we can not apply *pass-down* directly. On one hand, since we consider non-uniform capacity upper bounds, the inequality $\sum_{j=1}^{a}U_{u_j}\ge p$ may not be satisfied. We need to choose open centers carefully such that at least p vertices can be served. On the other hand, we can not open multiple facilities in a single vertex by hard capacities. Thus, we need the following lemma to modify the open center set F.

Lemma 13. Given an induced $(L,\{U_u\},\emptyset,p)$-Center instance induced by $G=(V=\mathcal{C}\cup\mathcal{F},E)$ where $|N_G(u)|\ge U_u$ for each $u\in\mathcal{F}$ and an open set $F=\{u_1,u_2,\cdots,u_a\}$ computed by *pass-up*, there exists a polynomial time algorithm that finds another open set $F'=\{u'_1,u'_2,\ldots,u'_a\}$ satisfying the following properties:

1. F' is a single set.
2. For all $1\le i\le a$, we have $\mathrm{d}_G(u_1,u'_1)\le 6$.
3. $\sum_{i=1}^{a}U_{u'_i}\ge p$.

We will prove the above lemma later. By Lemma 13, we are ready to prove Theorem 12.

Proof of Theorem 12. By Lemma 13, we obtain another open set $F'=\{u'_1,u'_2,\ldots,u'_a\}$. We first modify U_{u_i} to be $U_{u'_i}$ for all $1\le i\le a$. Then we apply the procedure *pass-down* according to the modified capacities. By Lemma 10, we obtain a distance-5 solution (C,F,ϕ). Since $\sum_{i=1}^{a}U_{u'_i}\ge p$, at least p vertices are served by ϕ. Finally, for each vertex $v\in C$ and $u_i\in F$ such that $\phi(v)=u_i$, we reassign v to $u'_i\in F'$, i.e., let $\phi(v)=u'_i$. By Lemma 13, we obtain a feasible solution for the given induced $(L,\{U_u\},\emptyset,p)$-Center instance. Since $\mathrm{d}(u_i,u'_i)\le 6$ $(1\le i\le a)$, the capacitated center value of our solution is at most $5+6=11$. Combining with Lemma 4, we finish the proof.

Now we only need to prove Lemma 13.

[¶] Recall that we may remove some facilities from $\mathcal{F}$ such that this assumption is satisfied. Thus, the set $\mathcal{F}$ may be a subset of V.

Proof. We construct an undirected weighted bipartite graph $B = (F, \mathcal{C}; E_B)$ as follows: $(u_j, v) \in E_B$ ($u_j \in F, v \in \mathcal{C}$) if and only if $\mathrm{d}(u_j, v) \le 6$ and this edge has weight U_v. We then find a maximum-weight maximum-matching M on this graph B. We only need to verify that F is perfectly matched in M and the total weight of M is at least p. Suppose each $u_i \in \mathcal{F}$ is matched to $v_i \in \mathcal{C}$, we finish the proof by letting $u'_i = v_i$.

Define $O \subseteq \mathcal{F} \subseteq \mathcal{C}$ to be the optimal open center set. By Hall's theorem, we first prove the existence of a matching M_1 in B satisfying that every vertex in O is matched. For any subset $O' \subseteq O$, we assume by contradiction that $|N_B(O')| \le |O'|-1$, where $N_B(O') = \cup_{v \in O'} N_B(v)$. Define C' to be the set of vertices served by O' in the optimal solution. By the capacity lower bound, we have $|C'| \ge L|O'|$. Recall that S is the unassigned set of size $b < L$ obtained by pass-up. Therefore, we have $|C' \setminus S| > L(|O'|-1) \ge L|N_B(O')|$. On the other hand, for each vertex $v \in C' \setminus S$, we have $\mathrm{d}_G(v, \phi(v)) \le 5$ by Lemma 9. Thus, we conclude that $\mathrm{d}_G(\phi(v), O') = \min_{w \in O'} \mathrm{d}_G(\phi(v), w) \le 6$, which implies all $\phi(v) \in N_B(O')$. Since each open center $u_i \in F$ only serves L vertices by Lemma 9, we have $|C' \setminus S| \le L|N_B(O')|$ which is a contradiction. So we prove the existence of M_1. Note that the total weight of M_1 is at least p.

Note that there exists a matching on B such that F is perfectly matched. We can achieve this property by matching each $u_i \in \mathcal{F}$ to an arbitrary vertex v_i such that $\phi(v_i) = u_i$. Then by Hungarian Algorithm, we can construct a matching M_2 by iteratively finding augmenting paths based on M_1, until all vertices in F are matched. Since any augmenting path can not make a matching vertex unmatched, we conclude that the total weight of M_2 is at least p. Thus, the maximum-weight maximum-matching M on B must satisfy that F is perfectly matched and the total weight is at least p.. □

(L,$\{U_u\}$,soft-$\emptyset$,p)-Supplier. Consider the (L,$\{U_u\}$,soft-$\emptyset$,p)-Supplier problem with soft capacities. Our technique is similar to (L,$\{U_u\}$,$\emptyset$,p)-Center except a difference procedure for choosing F' in Lemma 13. By Lemma 4, we again consider a given induced (L,$\{U_u\}$,soft-$\emptyset$,p)-Supplier instance induced by an undirected unweighted connected $G = (\mathcal{C} \cup \mathcal{F}, E)$. W.l.o.g., we assume that $|N_G(u)| \ge U_u$ for all facilities $u \in \mathcal{F}$. Similarly, we compute T, I, ξ and apply *pass-up* to compute $F = \{u_1, u_2, \ldots, u_a\}$, $S \subseteq \xi^{-1}(u^*)$ and ϕ. Before applying *pass-down*, we also need to find another open facility set $F' = \{u'_1, \ldots, u'_a\}$. Since we have soft capacities, we only need to require F' to satisfy Property 2 and 3 in Lemma 13. This is the reason why we can consider the supplier version. By the same technique as in the proof of Theorem 12, we have the following theorem.

Theorem 14. There exists a poly-time algorithm achieving approximation ratio 11 for (L,$\{U_u\}$,soft-$\emptyset$,p)-Supplier problem.

Proof. We only need to find another open facility set F' satisfying Property 2 and 3 in Lemma 13. In fact, we simply define $u'_i = \arg\max_{u \in \mathcal{F}: \mathrm{d}_G(u,u_i) \le 6} U_u$ for all $u_i \in F$. We only need to verify that $\sum_{i=1}^{a} U_{u'_i} \ge p$. Assume that the optimal open facility set is $O = \{u^*_1, u^*_2, \cdots, u^*_m\}$ and the optimal assignment function is ϕ^*. W.l.o.g., we assume that $U_{u^*_1} \ge U_{u^*_2} \ge \cdots \ge U_{u^*_m}$. We only need to find an injection $\sigma : \{1, 2, \cdots, w\} \to \{1, 2, \cdots, a\}$ such that $U_{u^*_j} \le U_{u'_{\sigma(j)}}$.

The injection can be found greedily. Suppose $\sigma(1), \sigma(2), \cdots, \sigma(j-1)$ have been decided for some $1 \le j \le w$, we want to decide $\sigma(j)$. Since each $u_i \in F$ only serves L clients by ϕ and the

unassigned client set S is of size $b < L$, we have the following property by counting:

$$A := \left(\bigcup_{i=1}^{j} \phi^{*-1}(u_i^*)\right) \setminus \left(S \cup \bigcup_{i=1}^{j-1} \phi^{-1}(u_{\sigma(i)})\right) \neq \emptyset.$$

Arbitrarily pick a client $v \in A$. Define $\sigma(j) = \phi(v)$. By the definition of O, there exists some $1 \leq j' \leq j$ such that $\mathrm{d}_G(v, u_{j'}^*) \leq 1$. Therefore, $\mathrm{d}_G(u_{\sigma(j)}, u_{j'}^*) \leq 6$. Thus, we have $U_{u_j^*} \leq U_{u_{j'}^*} \leq U_{u'_{\sigma(j)}}$ by the definition of $u'_{\sigma(j)}$. The proof is complete. □

3 Capacitated k-Center with Two-Sided Bounds and Outliers

In this section, we study the capacitated k-center problems with two-sided bounds, with or without outliers, and give approximation algorithms. We consider the case that all vertices have a uniform capacity lower bound $L_v = L$, while the capacity upper bounds can be either uniform or non-uniform. Our goal is to propose approximation algorithms with constant approximation ratio. Similar to (An *et al.* , 2015; Kociumaka & Cygan, 2014), we use the standard LP relaxation and the rounding procedure distance-r transfer. We will first extend the distance-r transfer procedure for two-sided bounds.

3.1 LP Formulation

We first give a natural LP relaxation for ($\{L_u\}$,$\{U_u\}$,k,p)-Supplier.

Definition 15. (Distance-r relaxation $\mathsf{LP}_r(G)$) Given an ($\{L_u\}$,$\{U_u\}$,k,p)-Supplier instance, the following feasibility $\mathsf{LP}_r(G)$ that fractionally verifies whether there exists a solution that assigns at least p clients to an open center of distance at most r:

$0 \leq x_{uv}, y_u \leq 1,$	$\forall u \in \mathcal{F}, v \in \mathcal{C};$
$x_{uv} = 0,$	if $\mathrm{d}(u,v) > r;$
$x_{uv} \leq y_u,$	$\forall u \in \mathcal{F}, v \in \mathcal{C};$
$\sum_{u \in \mathcal{F}} y_u = k;$	
$\sum_{u \in \mathcal{F}, v \in \mathcal{C}} x_{uv} \geq p;$	
$\sum_{u \in \mathcal{F}} x_{uv} \leq 1,$	$\forall v \in \mathcal{C};$
$L_u y_u \leq \sum_{v \in \mathcal{C}} x_{uv} \leq U_u y_u,$	$\forall u \in \mathcal{F}.$

Here x_{uv} is called an assignment variable representing the fractional amount of assignment from client v to center u, and y_u is called the opening variable of $u \in \mathcal{F}$. For convenience, we use x, y to represent $\{x_{uv}\}_{u \in \mathcal{F}, v \in \mathcal{C}}$ and $\{y_u\}_{u \in \mathcal{F}}$, respectively.

By Definition 3, $\mathsf{LP}_1(G)$ must have a feasible solution for any induced($\{L_u\}$,$\{U_u\}$,k,p)-Supplier instance. Assume that we have a feasible fractional solution (x, y) of $\mathsf{LP}_1(G)$. We want to obtain a distance-ρ solution by rounding (x, y). We then recall a rounding procedure called distance-r transfer.

3.2 Distance-r Transfer

We first extend the definition of distance-r transfer proposed in (An *et al.* , 2015; Kociumaka & Cygan, 2014) by adding the third condition. For a vertex $a \in \mathcal{C} \cup \mathcal{F}$ and a set $B \subseteq \mathcal{C} \cup \mathcal{F}$, we define $\mathrm{d}(a, B) = \min_{b \in B} \mathrm{d}(a, b)$.

Definition 16. Given an $(\{L_u\},\{U_u\},k,p)$-Supplier instance and $y \in \mathbb{R}^{\mathcal{F}}_{\geq 0}$, a vector $y' \in \mathbb{R}^{\mathcal{F}}_{\geq 0}$ is a distance-r transfer of y if

1. $\sum_{u \in \mathcal{F}} y'_u = \sum_{u \in \mathcal{F}} y_u$;
2. $\sum_{w \in \mathcal{F}: \mathrm{d}(w,W) \leq r} U_w y'_w \geq \sum_{u \in W} U_u y_u$ for all $W \subseteq \mathcal{F}$;
3. $\sum_{w \in \mathcal{F}: \mathrm{d}(w,W) \leq r} L_w y_w \geq \sum_{u \in W} L_u y'_u$ for all $W \subseteq \mathcal{F}$.

If y' is a characteristic vector of $F \subseteq \mathcal{F}$, we say that F is an integral distance-r transfer of y.

Recall that the first condition says that a transfer should not change the total number of open centers. By an argument using Hall's theorem as in (An *et al.* , 2015; Kociumaka & Cygan, 2014), the second condition is important for satisfying the capacity upper bounds. In this paper, we add the third condition to satisfy the capacity lower bounds. Like in (An *et al.* , 2015; Kociumaka & Cygan, 2014), an integral distance-r transfer of the fractional solution of $\mathsf{LP}_r(G)$ already gives a distance-$(r+1)$ solution by the following lemma.

Lemma 17. Given an $(\{L_u\},\{U_u\},k,p)$-Supplier problem, assume (x, y) is a feasible solution of $\mathsf{LP}_1(G)$ and $F \subseteq \mathcal{F}$ is an integral distance-r transfer of y. Then one can find a distance-$(r+1)$ solution (C, F, ϕ) in polynomial time.

Proof. Consider a bipartite graph $H = (\mathcal{C}, \mathcal{F}, E_H)$ with $(v, u) \in E_H$ $(v \in \mathcal{C}, u \in \mathcal{F})$ if $\mathrm{d}_G(v, u) \leq r + 1$. Modify H to obtain $H_1 = (\mathcal{C}, F_1, E_{H_1})$ by removing vertices from $\mathcal{F} \setminus F$ and duplicating each vertex $u \in F$ to its capacity lower bound, i.e. L_u times. Then we show that there exists a matching M_1 such that every vertex in F_1 is matched. By Hall's theorem, we need to prove that $|\{v \in \mathcal{C} : \mathrm{d}_{H'}(v, W) \leq r + 1\}| \geq |W|$ for any $W \subseteq F_1$. By the construction of F_1, we in fact only need to prove that for each $W \subseteq F$, we have the following:

$$|\{v \in \mathcal{C} : \mathrm{d}_G(v, W) \leq r + 1\}| \geq \sum_{u \in W} L_u.$$

By the third condition of Definition 16, we know $\sum_{u \in W} L_u \leq \sum_{w \in \mathcal{F}: \mathrm{d}_G(w,W) \leq r} L_w y_w$. Then by the LP constraint, we have the following inequality:

$$\begin{aligned} \sum_{u \in W} L_u &\leq \sum_{w \in \mathcal{F}: \mathrm{d}_G(w,W) \leq r} L_w y_w \leq \sum_{w \in \mathcal{F}: \mathrm{d}_G(w,W) \leq r} \sum_{v \in \mathcal{C}: \mathrm{d}_G(w,v) \leq 1} x_{wv} \\ &\leq \sum_{v \in \mathcal{C}: \mathrm{d}_G(v,W) \leq r+1} \sum_{w \in \mathcal{F}} x_{wv} \leq \sum_{v \in \mathcal{C}: \mathrm{d}_G(v,W) \leq r+1} \sum_{w \in \mathcal{F}} 1 = |\{v \in \mathcal{C} : \mathrm{d}_G(v, W) \leq r + 1\}|. \end{aligned}$$

We then modify H to obtain another bipartite graph $H_2 = (\mathcal{C}, F_2, E_{H_2})$ by removing vertices from $\mathcal{F} \setminus F$ and duplicating each vertex $u \in F$ to its capacity upper bound, i.e. U_u times. Since

$L_u \leq U_u$ for all $u \in \mathcal{F}$, we can consider F_1 as a subset of F_2. Moreover, we can consider M_1 as a matching of H_2 satisfying that for each $u \in F$, there are exactly L_u duplicates of u to be matched. On the other hand, there exists a matching M_2 of H_2 satisfying that there are at least p clients $v \in \mathcal{C}$ that are matched (Kociumaka & Cygan, 2014, Lemma 9). Finally, we show that it is possible to modify M_1 to be a matching with the same property as M_2. Starting with M_1, we iteratively find an augmenting path on H_2 such that there is one more client that are matched. We stop until there are at least p matched clients. The feasibility is not hard by Hungarian Algorithm. Since any augmenting path can not make a matched vertex unmatched, we conclude that for each $u \in F$ there are at least L_u matched duplicates. This matching can be found in polynomial time, and $|F| = k$ follows from the first condition of Definition 16. □

Note that the capacity lower bound is uniform, we have the following lemma for distance-r transfer.

Lemma 18. Given an (L,$\{U_u\}$,k,p)-Supplier instance and $y \in \mathbb{R}^{\mathcal{F}}_{\geq 0}$, assume that a vector $y' \in \mathbb{R}^{\mathcal{F}}_{\geq 0}$ satisfies the following conditions:

1. $\sum_{u\in\mathcal{F}} y'_u = \sum_{u\in\mathcal{F}} y_u$;
2. $\sum_{w\in\mathcal{F}:\mathrm{d}(w,W)\leq r} U_w y'_w \geq \sum_{u\in W} U_u y_u$ for all $W \subseteq \mathcal{F}$;
3. There exists a function $g : \mathcal{F} \times \mathcal{F} \to \mathbb{R}_{\geq 0}$ satisfying the following conditions:
   (a) for all $u, w \in \mathcal{F}$, if $\mathrm{d}(u, w) > r$, we have $g(u, w) = 0$;
   (b) for all $u \in \mathcal{F}, y_u = \sum_{w\in\mathcal{F}} g(u, w), y'_u = \sum_{w\in\mathcal{F}} g(w, u)$.

Then y' is a distance-r transfer of y. We call this y' a local distance-r transfer of y.

Proof. By Definition 16, we only need to verify that $\sum_{w\in\mathcal{F}:\mathrm{d}(w,W)\leq r} L_w y_w \geq \sum_{u\in W} L_u y'_u$ for all $W \subseteq \mathcal{F}$. Since the capacity lower bound is uniform, we can simplify the above inequality as follows: $\sum_{w\in\mathcal{F}:\mathrm{d}(w,W)\leq r} y_w \geq \sum_{u\in W} y'_u$. We finish the proof by the following argument:

$$\begin{aligned}\sum_{u\in W} y'_u &= \sum_{u\in W}\sum_{w\in\mathcal{F}} g(w,u) = \sum_{u\in W}\sum_{w\in\mathcal{F}:\mathrm{d}(w,u)\leq r} g(w,u) \leq \sum_{u\in W}\sum_{w\in\mathcal{F}:\mathrm{d}(w,W)\leq r} g(w,u)\\ &= \sum_{w\in\mathcal{F}:\mathrm{d}(w,W)\leq r}\sum_{u\in W} g(w,u) = \sum_{w\in\mathcal{F}:\mathrm{d}(w,W)\leq r} y_w.\end{aligned}$$

The second equality follows from condition 3a. □

The function g in the above lemma can be considered as a flow from the facility set $\mathcal{F}$ to its duplicate $\mathcal{F}'$, with the property that no flow enters $w \in \mathcal{F}'$ from any facility $u \in \mathcal{F}$ such that $\mathrm{d}(u, w) > r$. In other words, such a distance-r transfer does not make a transfer of opening variables from some facility $u \in \mathcal{F}$ to a 'too far' facility $w \in \mathcal{F}$ satisfying that $\mathrm{d}(u, w) > r$.

Similarly, if the capacity upper bound is uniform, we have the following lemma which may give some intuition for solving the ($\{L_u\}$,U,k,p)-Supplier problems in the future.

Lemma 19. Given an ($\{L_u\}$,U,k,p)-Supplier instance and $y \in \mathbb{R}^{\mathcal{F}}_{\geq 0}$, assume that a vector $y' \in \mathbb{R}^{\mathcal{F}}_{\geq 0}$ satisfies the following conditions:

1. $\sum_{u\in\mathcal{F}} y'_u = \sum_{u\in\mathcal{F}} y_u$;
2. $\sum_{w\in\mathcal{F}:\mathrm{d}(w,W)\le r} L_w y_w \ge \sum_{u\in W} L_u y'_u$ for all $W \subseteq \mathcal{F}$
3. There exists a function $g : \mathcal{F}\times\mathcal{F} \to \mathbb{R}_{\ge 0}$ satisfying the following conditions:
   (a) $\forall u, w \in \mathcal{F}$, if $\mathrm{d}(u, w) > r$, we have $g(u, w) = 0$;
   (b) $\forall u \in \mathcal{F}, y_u = \sum_{w\in\mathcal{F}} g(u, w), y'_u = \sum_{w\in\mathcal{F}} g(w, u)$.

Then y' is a distance-r transfer of y.

3.3 Capacitated k-Center with Two-Sided Bounds and Outliers

Now we are ready to solve the (L,$\{U_u\}$,k,p)-Supplier problem. By Lemma 17 and 18, we only need to find an integral local distance-r transfer. In (An *et al.* , 2015; Kociumaka & Cygan, 2014), there exists a polynomial time algorithm that finds an integral distance-r transfer for the (L,$\{U_u\}$,k,p)-Supplier problem with capacity lower bound $L = 0$. Their algorithm can be naturally generalized to the case that $L > 0$. Thus, we only need to show that their rounding scheme also finds a local distance-r transfer. In this section, we briefly state the rounding schemes for the k-center problem with only capacity upper bounds in previous work (An *et al.* , 2015; Kociumaka & Cygan, 2014), and show that they find integral local distance-r transfers. We first have the following observation.

Observation 3.1. Given an (L,$\{U_u\}$,k,p)-Supplier instance, let $y, y', y'' \in \mathbb{R}^{\mathcal{F}}_{\ge 0}$. Assume y' is a local distance-r transfer of y and y'' is a local distance-r' transfer of y'. Then y'' is a local distance-$(r + r')$ transfer.

Though the authors consider many variants in (An *et al.* , 2015; Kociumaka & Cygan, 2014), the rounding schemes are quite similar. We first take the (0,$\{U_u\}$,k)-Center problem as an example. An *et al.* (2015) proposed a rounding scheme for this problem by reducing to a tree instance defined as follows.

Definition 20. (Tree instance) A tree instance is defined as a tuple $(T, \{L_v\}, \{U_v\}, y)$, where $T = (V, E)$ is a rooted tree with a capacity interval $[L_v, U_v]$ for each facility $v \in V$ ($L_v, U_v \in \mathbb{Z}_{\ge 0}$ and $L_v \le U_v$), a distance function $\mathrm{d}_T : V \times V \to \mathbb{R}_{\ge 0}$ induced by T, and opening variables $y \in [0, 1]^V$ satisfying that $\sum_{u\in V} y_v$ is an integer and $y_v = 1$ for every non-leaf facility $v \in V$.

The rounding procedure is as follows.

1. Find a fractional solution (x, y) of the standard LP.
2. Construct a distance-1 transfer y' of y by aggregating one opening unit to each vertex in a set V from its neighbors. [‖] Based on y', construct a tree instance $(T = (V, E_T), 0, \{U_u\}, y')$.
3. Find an integral distance-2 transfer y'' of y' based on the tree induced distance by (An *et al.* , 2015, Lemma 9). Since the tree T satisfies that the original distance between any two adjacent vertices is at most 3, y'' is also a distance-6 transfer of y' based on the original distance.

[‖] Here, V is a set of auxiliary vertices augmenting to the original instance. See (An *et al.* , 2015) for more details.

4. Finally, find an open set $F \in \mathcal{F}$ which is also a distance-1 transfer of y''. By the transitivity of distance-r transfer, F is a distance-8 transfer of y. By Lemma 17, F is a distance-9 solution.

Our goal is to show that the above rounding scheme can be generalized to the case that $L > 0$. We use exactly the same rounding scheme except that the initial fractional solution (x, y) should satisfy all capacity lower bounds. We then find an open set F as an integral distance-8 transfer of y, and we want to show that F is also local. It is not hard to verify the local property is satisfied in Step 2 and Step 4. Thus, we only need to check Step 2 for the tree instance. In fact, we have the following lemma.

Lemma 21. Given a tree instance $(T, L, \{U_u\}, y)$, one can find in polynomial time an integral local distance-2 transfer of y.

Proof. We recall the rounding scheme of tree instances in (An *et al.* , 2015).

If T only consists of at most 1 vertex, then y is an integral vector itself, yielding an integral local distance-1 transfer of y.

If T consists of more than 1 vertex, we construct a distance-2 transfer by a recursion. We first find a non-leaf vertex $r \in T$ whose children are all leaves. We use $T_1 = (V_1, E_1)$ to denote the subtree rooted at r, and use $T_2 = (V_2, E_2)$ to denote the subtree $(T \setminus T_1) \cup \{r\}$ by deleting all children of r from T. Define $y^1 \in [0,1]^{V_1}$ and $y^2 \in [0,1]^{V_2}$ as follows: $y^1|_{V_1 \setminus \{r\}} = y|_{V_1 \setminus \{r\}}$, $y^2|_{V_2 \setminus \{r\}} = y|_{V_2 \setminus \{r\}}$ and $y_r = y^1_r + y^2_r$ where $0 \le y^1_r < 1$ and $\sum_{u \in V_1} y^1_u$ is an integer. Since $y_r = 1$, we can always find such values y^1_r and y^2_r. Define $Y = \sum_{u \in V_1} y^1_u$. By the above definition, we observe that Y is an integer in $[0, |V_1| - 1]$. We use F_1 to denote the collection of Y vertices among the children of r of highest capacity upper bounds. We use u^* to denote the child of r of the Y-th highest capacity upper bound. Note that $u^* \in F_1$. We consider the following cases.

1. If $U_r < U_{u^*}$, then F_1 is an integral distance-2 transfer of y_1 by (An *et al.* , 2015, Lemma 9). To prove that F_1 is also a local distance-2 transfer of y_1, we only need to verify the third condition in Lemma 18. Note that $|F_1| = Y$. There must exist a function $g : V_1 \times V_1 \to \mathbb{R}_{\ge 0}$ satisfying condition 3b. On the other hand, the diameter of the tree T_1 is at most 2. So the function g must satisfy condition 3a, which finishes the proof. Then the algorithm recursively solves a smaller tree instance and obtains $F_2 \subseteq V_2$ as an integral local distance-2 transfer of y_2. Since $F_1 \cap F_2 = \emptyset$, we conclude that $F := F_1 \cup F_2$ is an integral local distance-2 transfer of y.

2. If $U_r = U_{u^*}$, then both F_1 and $F'_1 := (F_1 \cup \{r\}) \setminus \{u^*\}$ are integral local distance-2 transfers of y_1 by the same argument as in the first case. In this case, we first recursively solves a smaller tree instance and obtains $F_2 \subseteq V_2$ as an integral local distance-2 transfer of y_2. Then we define
$$F := \begin{cases} F_1 \cup F_2, & r \in F_2 \\ F'_1 \cup F_2, & r \notin F_2 \end{cases},$$
which is an integral local distance-2 transfer of y.

3. If $U_r > U_{u^*}$, we modify the capacity upper bound U_r by $U_r \leftarrow U_{u^*}$ and reduce to the second case. In fact, we can similify the argument of this case in (An *et al.* , 2015). We only need to verify the second condition in Definition 16. By the seond case, $y'_r = 1$ always holds, i.e. $r \in F$. Since F is an integral distance-2 transfer of y by replacing U_r by a smaller value U_{u^*},

we can check that the second condition in Definition 16 still holds for the original capacity U_r.

□

Thus, we directly have the following theorem by An *et al.* (2015).

Theorem 22. There is a polynomial time 9-approximation algorithm for the $(L,\{U_u\},k)$-Center problem. For the uniform capacity upper bound version, the (L,U,k)-Center problem admits a 6-approximation.

Theorem 23. There is a polynomial time 13-approximation algorithm for the $(L,\{U_u\},k)$-Supplier problem. For the uniform capacity upper bound version, the (L,U,k)-Supplier problem admits a 9-approximation.

The case with outliers. For the $(0,\{U_u\},k,p)$-Supplier problem with outliers, Kociumaka & Cygan (2014) gave an algorithm to find an integral distance-24 transfer. W.l.o.g., they assume the optimal capacitated k-supplier value is 1. For each facility $u \in \mathcal{F}$, define $deg(u) = |\{v \in \mathcal{C}, \mathrm{d}(u,v) \leq 1\}|$. They need to assume that $U_u \leq deg(u)$ by setting $U_u \leftarrow \min\{U_u, deg(u)\}$. The same as in Section 2, we can still make this assumption. Another difference is that they reduce the problem to a tree instance via a data structure called *skeleton*, see (Kociumaka & Cygan, 2014, Definition 3). We need to verify the existence of a skeleton for the case that $L > 0$. Fortunately, we can first modify $L = 0$ and use the algorithm in (Kociumaka & Cygan, 2014) to find a skeleton S. It is not hard to verify that S is also a skeleton for L. Similarly, by (Kociumaka & Cygan, 2014), we have the following theorem.

Theorem 24. There is a polynomial time 25-approximation algorithm for the $(L,\{U_u\},k,p)$-Supplier problem and the $(L,\{U_u\},\text{soft-}k,p)$-Supplier problem. For the uniform capacity upper bound version, the (L,U,k,p)-Supplier problem admits a 23-approximation, and the $(L,U,\text{soft-}k,p)$-Supplier problem admits a 13-approximation.

Note that taking $\mathcal{C} = \mathcal{F} = V$ shows that the k-supplier problem generalizes the k-center problem, and consequently gives the same approximation bounds for the latter.

Corollary 25. There is a polynomial time 25-approximation algorithm for the $(L,\{U_u\},k,p)$-Center problem and the $(L,\{U_u\},\text{soft-}k,p)$-Center problem. For the uniform capacity upper bound version, the (L,U,k,p)-Center problem admits a 23-approximation, and the $(L,U,\text{soft-}k,p)$-Center problem admits a 13-approximation.

References

Aggarwal, Gagan, Panigrahy, Rina, Feder, Tomás, Thomas, Dilys, Kenthapadi, Krishnaram, Khuller, Samir, & Zhu, An. 2010. Achieving anonymity via clustering. *ACM Trans. Algorithms*, **6**(3), 49:1–49:19.

Ahmadian, Sara, & Swamy, Chaitanya. 2016. Approximation Algorithms for Clustering Problems with Lower Bounds and Outliers. *arXiv preprint arXiv:1608.01700*.

An, Hyung-Chan, Bhaskara, Aditya, Chekuri, Chandra, Gupta, Shalmoli, Madan, Vivek, & Svensson, Ola. 2015. Centrality of trees for capacitated k-center. *Mathematical Programming*, **154**(1-2), 29–53.

Arya, Vijay, Garg, Naveen, Khandekar, Rohit, Meyerson, Adam, Munagala, Kamesh, & Pandit, Vinayaka. 2004. Local search heuristics for k-median and facility location problems. *SIAM Journal on computing*, **33**(3), 544–562.

Barilan, Judit, Kortsarz, Guy, & Peleg, David. 1993. How to allocate network centers. *Journal of Algorithms*, **15**(3), 385–415.

Charikar, Moses, & Guha, Sudipto. 2005. Improved combinatorial algorithms for facility location problems. *SIAM Journal on Computing*, **34**(4), 803–824.

Charikar, Moses, Guha, Sudipto, Tardos, Éva, & Shmoys, David B. 1999. A constant-factor approximation algorithm for the k-median problem. *Pages 1–10 of: Proceedings of the thirty-first annual ACM symposium on Theory of computing*. ACM.

Charikar, Moses, Khuller, Samir, Mount, David M, & Narasimhan, Giri. 2001. Algorithms for facility location problems with outliers. *Pages 642–651 of: Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*. Society for Industrial and Applied Mathematics.

Chen, Danny Z, Li, Jian, Liang, Hongyu, & Wang, Haitao. 2016. Matroid and knapsack center problems. *Algorithmica*, **75**(1), 27–52.

Cygan, Marek, Hajiaghayi, MohammadTaghi, & Khuller, Samir. 2012. LP Rounding for k-Centers with Non-uniform Hard Capacities. *Pages 273–282 of: Proceedings of the 2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*. IEEE Computer Society.

Dick, Travis, Li, Mu, Pillutla, Venkata Krishna, White, Colin, Balcan, Maria Florina, & Smola, Alex. 2015. Data Driven Resource Allocation for Distributed Learning. *arXiv preprint arXiv:1512.04848*.

Ene, Alina, Har-Peled, Sariel, & Raichel, Benjamin. 2013. Fast clustering with lower bounds: No customer too far, no shop too small. *arXiv preprint arXiv:1304.7318*.

Fernandes, Cristina G, de Paula, Samuel P, & Pedrosa, Lehilton LC. 2016. Improved Approximation Algorithms for Capacitated Fault-Tolerant k-Center. *Pages 441–453 of: Latin American Symposium on Theoretical Informatics*. Springer.

Gonzalez, Teofilo F. 1985. Clustering to minimize the maximum intercluster distance. *Theoretical Computer Science*, **38**, 293–306.

Guha, Sudipto, & Khuller, Samir. 1999. Greedy strikes back: Improved facility location algorithms. *Journal of algorithms*, **31**(1), 228–248.

Hochbaum, Dorit S, & Shmoys, David B. 1985. A best possible heuristic for the k-center problem. *Mathematics of operations research*, **10**(2), 180–184.

Jain, Kamal, & Vazirani, Vijay V. 2001. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and Lagrangian relaxation. *Journal of the ACM (JACM)*, **48**(2), 274–296.

Jain, Kamal, Mahdian, Mohammad, & Saberi, Amin. 2002. A new greedy approach for facility location problems. *Pages 731–740 of: Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*. ACM.

Khuller, Samir, & Sussmann, Yoram J. 2000. The capacitated k-center problem. *SIAM Journal on Discrete Mathematics*, **13**(3), 403–418.

Kociumaka, Tomasz, & Cygan, Marek. 2014. Constant factor approximation for capacitated k-center with outliers. *arXiv preprint arXiv:1401.2874*.

Korupolu, Madhukar R, Plaxton, C Greg, & Rajaraman, Rajmohan. 2000. Analysis of a local search heuristic for facility location problems. *Journal of algorithms*, **37**(1), 146–188.

Li, Jian, Yi, Ke, & Zhang, Qin. 2010. Clustering with diversity. *Pages 188–200 of: International Colloquium on Automata, Languages, and Programming*. Springer.

Li, Shi. 2013. A 1.488 approximation algorithm for the uncapacitated facility location problem. *Information and Computation*, **222**, 45–58.

Li, Shi, & Svensson, Ola. 2016. Approximating k-median via pseudo-approximation. *SIAM Journal on Computing*, **45**(2), 530–547.

Sweeney, Latanya. 2002. k-anonymity: A model for protecting privacy. *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems*, **10**(05), 557–570.

A Proof of Lemma 4

Proof. Suppose we have such a polynomial time algorithm A as described in the lemma, and the optimal capacitated k-supplier value of the input $(\{L_u\},\{U_u\},k,p)$-Supplier instance is $r > 0$. We want to give a polynomial time algorithm that outputs a feasible solution with a capacitated k-supplier value at most ρr.

We first guess the value of r. Note that there are at most $|\mathcal{C}| \times \mathcal{F}|$ possible values for the optimal capacitated k-supplier value r, we can enumerate all of them. Assume that we know the value r. We connect each pair (u, v) ($u \in \mathcal{F}$, $v \in \mathcal{C}$) if and only if $\mathrm{d}(u, v) \leq r$, and obtain an undirected graph $G = (\mathcal{F} \cup \mathcal{C}; E)$. Suppose G can be decomposed into different connected components $G_i = (\mathcal{C}_i \cup \mathcal{F}_i, E_i)$ ($i = 1, 2, \cdots, m$). Since the optimal capacitated k-supplier value of the input $(\{L_u\},\{U_u\},k,p)$-Supplier instance is r, there must exist client sets $C_i \subseteq \mathcal{C}$ of size

$p_i \in \mathbb{Z}_{\geq 0}$ satisfying that $\sum_{i=1}^{m} p_i \geq p$, open facility sets $F_i \subset \mathcal{F}_i$ of size k_i satisfying that $\sum_{i=1}^{m} k_i = k$, and an assignment of every client $v \in C_i$ $(i = 1, 2, \cdots, m)$ to an open facility $u \in F_i$ such that $\mathrm{d}(u, v) \leq r$ and the capacity constraints are satisfied.

For each connected component G_i, consider the distance function $\mathrm{d}_{G_i} : (\mathcal{C}_i \cup \mathcal{F}_i) \times (\mathcal{C}_i \cup \mathcal{F}_i) \to \mathbb{R}_{\geq 0}$ induced by G_i. Combining with the above argument, we have that the optimal capacitated k_i-supplier value of the induced $(\{L_u\},\{U_u\},k_i,p_i)$-Supplier instance induced by G_i is at most 1. If we have the exact value of k_i and p_i, we can find a feasible solution with a capacitated k_i-supplier value at most ρ for this induced instance by the assumption in the lemma. Observe that for any $s_1, s_2 \in \mathcal{C}_i \cup \mathcal{F}_i$, we have that the original distance $\mathrm{d}(s_1, s_2)$ between s_1 and s_2 satisfies that $\mathrm{d}(s_1, s_2) \leq r \cdot \mathrm{d}_{G_i}(s_1, s_2)$. Thus, the union of all such feasible solutions of each G_i is a ρ-approximation feasible solution of the original $(\{L_u\},\{U_u\},k,p)$-Supplier problem. We remain to show that we can find values k_i, p_i in polynomial time.

Let the boolean value $A_i[k'][p']$ indicate whether the algorithm A finds a feasible solution with capacitated k'-supplier value at most ρ for the induced $(\{L_u\},\{U_u\},k_i,p_i)$-Supplier instance induced by G_i. Let the boolean value $K[i][k'][p']$ indicate whether there exists a feasible solution that opens k' facilities in $\mathcal{F}_1 \cup \mathcal{F}_2 \cup \cdots \cup \mathcal{F}_i$ to cover at least p' clients in $\mathcal{C}_1 \cup \mathcal{C}_2 \cup \cdots \cup \mathcal{C}_i$, satisfying all capacity constraints and the capacitated k'-supplier value (the original $(\{L_u\},\{U_u\},k,p)$-Supplier instance) is at most ρr. Initially, let $K[1][k'][p'] = A_1[k'][p']$. For $i > 1$, we update

$$K[i][k'][p'] = \bigvee_{\substack{0 \leq k^* \leq k' \\ 0 \leq p^* \leq p'}} (K[i-1][k'-k^*][p'-p^*] \wedge A_i[k^*][p^*]).$$

Note that $K[m][k][p] = \mathsf{true}$ since all $A_i[k_i][p_i] = \mathsf{true}$. Thus in polynomial time we can find values k_i and p_i $(i = 1, 2, \cdots, m)$ such that $A_i[k_i][p_i] = \mathsf{true}$ and $\sum_{i=1}^{m} k_i = k, \sum_{i=1}^{m} p_i \geq p$. □